How to choose? - An Introduction to Decision Theory

0. Introduction

Charles Darwin once described the mathematician *'as a blind man in a dark room looking for a black cat which isn't there'*, whilst Ellis chose to go with the desciption as one who *'has reached the highest rung on the ladder of human thought'*. Whichever analogy you feel is more suited, I hope that, at the very least, some basics can be agreed upon. In my limited experience, and to take a hugely sweeping generalisation, the mathematician appears to be one who observes and models the current state of the world in an attempt to explain it, and, in the most ideal of circumstances, predict its future. The areas I have chosen to explore in this essay - preference, utility and their applications in decision making - pose no major exception to this.

Over the course of this essay I hope to introduce the reader to some of the key concepts and ideas underpinning how choices can be made mathematically and including some, hopefully stimulating, examples of how the mathematical theory behind choice can be applied to real-life situations to help put this into context.

One of the most interesting aspects of choice is that it is largely subjective [1]. It has been suggested that situations in which mathematical decision making can applied fall into 3 categories; *'Under Certainty'*, *'With Risk'* and *'Under Strict Uncertainty'*. The theory surrounding the case of *'Under Strict Uncertainty'* has a tendency to be rather algorithmic in nature and is plagued by conflicting results stemming from a number of competing schools of thought [2]. Consequently, during the course of this essay I intend to concentrate on the cases of decision making *'Under Certainty'* - wherein the environment is known and, as such, consequence can be predicted - and *'With Risk'* - wherein probabilities can be attached to the occurrence of each possible environment.

As it is human nature to retrospectively obsess over choices, motivation should be rather clear. Given a choice between a number of options the ability to consistantly guarantee that the optimal choice will be made is a invaluable skill regardless of the industry in which it is applied. I conjecture that Businesses which allocate their budget based upon decision mathematics and optimisation algorithms will almost always outperform any competitor who has done so based upon instinct or any other method you might care to name. Likewise, I believe that individuals investing their money in this way will profit more consistantly, and to a greater extent, than those who have not.

[1] Dependent on who is making the choice, and the environment in which it is being made
[2] See, for example, the Max-Min and Min-Max Choice Rules

1. The Nature of Preference

Preference is defined, literally, as *'The selecting of someone or something over another or others'* which should, at the very least, hint at how closely related preference and choice are. Although, linguistically speaking, the concepts are fairly synomynous, the connection is not quite as definitive as you might hope when considered in a mathematical sense. Indeed, although the two ideas are still indisputably related (as will be made clear later) preference is actually a substancially weaker condition and, consequently, is used to form a foundation from which theory encompassing choice can be constructed.

Preference itself is defined using a specialised type of *Binary Relation* known, rather uncreatively, as a *Preference Relation.* A *Binary Relation* is defined formally in the following way:

1.1 Definition: Binary Relation *(aka dyadic or 2-place relation)*
A *Binary Relation* is defined as the ordered triple (A, B, G) where A and B are arbitrary sets (or classes) and $G \subset A \times B$.

Notation:
1. If $a \in A$ is related to $b \in B$ by R then aRb.
2. When $A = B$ it is said that R is a *Binary Relation on A*

It should be noted that the fact that a *Binary Relation* is, in part, defined as a subset of the cartesian product is incredibly important. Not only does it enable the application of known set operations to the subset of the cartesian product yielded by any given *Binary Relation*[1] , but it also means that a model for choice with the desirable property that a set of optimal choices (as opposed to just one) can be produced.

There are a huge number of properties that can be attached to a *Binary Relation* to describe, in a more in depth fashion, how it actually acts. The vast majority of these are rather irrelevant to the process of defining a *Preference Relation* and so have been omitted[2], but those that are immediately relevant follow:

1.2 Definition: Complete
A *Binary Relation* is said to be *Complete* if $\forall x, y \in A$; xRy or yRx (but not both).

1.3 Definition: Asymmetric
A *Binary Relation* is said to be *Asymmetric* if $\forall x, y \in A$, $xRy \Rightarrow yR^c x$

[1]See: http://en.wikipedia.org/wiki/Set_(mathematics)♯Basic_ operations
[2]See [1]. Page 19

1.4 Definition: Reflexive
A *Binary Relation* is said to be *Reflexive* if $\forall x \in A$, xRx

These properties can be summarised contextually as follows:
Completeness: Given two elements x and y, a comparison can be made.
Asymmetry: If x is (strictly) prefered to y, then y isn't (strictly) prefered to x.
Reflexivity: Where given a choice between x and x, x is chosen.

Using these definitions, a *Preference Relation* can now be defined - this is done in the following way:

1.5.1 Definition: (Strict) Preference Relation
A *Preference Relation* P is an asymmetric *Binary Relation* on a set of possible actions A such that: For any two elements $a, b \in A$, the ordered pair (a, b) lies within the set of preferences P if a is (strictly) prefered to b.

The keen-eyed reader may have spotted that only one of the 3 previously considered properties *(Asymmetry/Completeness/Reflexivity (1.2 - 1.4))* was required to properly define a *(Strict) Preference Relation.* It may come as no surprise that there is actually an alternate definition utilising the remaining properties and based around the idea of *Weak Preference* (as opposed to a *Strict Preference*). This is written formally in the following manner:

1.5.2 Definition: (Weak) Preference Relation
A *Preference Relation*, P, is a complete, reflexive binary relation on a set of possible actions, A, such that: For any two elements $a, b \in A$ the ordered pair (a, b) lies within the set of preferences, P, if b is not prefered to a. *(i.e. a is prefered to, or considered indifferent from, b)*

When faced with two conflicting definitions. the obvious question that arises is: 'Which one is better suited to the task at hand?'. Whilst using *1.5.1* would mean that there are fewer ordered-pairs to consider overall, *1.5.2* has the advantage that it provides a far more intuitive way of dealing with "self-comparison". Fortuitously, it turns out that this is not a decision that needs to be made as a *Strict Preference Relation* and a *Weak Preference Relation* are indistinguishable[3].

1.6 Lemma *(Equivalence of Definitions)*
An asymmetric binary operation formalising strict preference *(See 1.5.1)* modelling a given collection of preferences, and a complete binary operation formalising weak preference *(See 1.5.2)* modelling the same collection of preferences are equivalent.

[3]When defined as in 1.5.1 and 1.5.2 (Note: Some definitions omit completeness from 1.5.2 and it is an essential property for the definitions to be equivalent)

Proof:
Suppose that, given some $x, y \in A$ and some complete binary operation formalising weak preference P, $(x, y) \in P$. Completeness implies that $(y, x) \notin P$ so $(y, x) \in P^c$ and asymmetry must hold. Additionally, completeness, by its definition, ensures strict preference. (If x and y are indifferent then either $(x, y) \in P$ or $(y, x) \in P$. Either way this contradicts x and y being indifferent). Hence having a complete binary relation formalising weak preference is the same as having an asymmetric binary operation formalising strict preference.

Now suppose that, given some $x, y \in A$ and some asymmetric binary operation formalising strict preference P $(x, y) \in P$. Asymmetry implies that $(y, x) \in P^c$ and so, given any $x, y \in A$, either xPy or yPx (Completeness). As strict preference is a stronger condition than weak preference, this is enough to show that having an asymmetric binary operation formalising strict preference is the same as having a complete binary relation formalising weak preference.

Hence the two definitions are equivalent.
—

Notation:
If a is *(Strictly) Prefered* to b then $a \succ b$
If a is considered to be *indifferent* from b then $a \sim b$
If a is considered to be *'at least as good as'* b then $a \succeq b$

A careful look at the notation above might draw attention to the similarity between $\succeq$ and $\geq$. Indeed, much like $\geq$ is the combination of the two symbols $>$ and $=$, it is true that $\succeq$ is the combination of $\succ$ *((Strictly) Prefered)*, and $\sim$ *(indifferent to)*. Continuing along these lines, it should be noted that the word "indifferent" is fairly synonymous with the word "equivalent" when discussing preferences.

From the definition, it can be seen that the basic idea behind a *Preference Relation* is to take any two elements from our set of possible actions and produce an ordered-pair that denotes preference clearly. What we are yet to show, however, is exactly how an element of the set of preferences interacts with any other element of the set. Although it may seem rather intuitive that the property of transitivity should hold when dealing with a concept as humanistic as preference, this is not necessarily the case as shown in the following example:

1.7 Example *(jan-ken-pon)*
Two players are participating in a of jan-ken-pon[4]. In case the concept is unfamiliar, in this game each player chooses one of 3 options: Rock (r), Paper (p), or Scissors (s) with the relationships that: Paper is prefered to Rock, Scissors to Paper and Rock to Scissors (i.e. The set of actions is $X = \{r, p, s\}$ and the

[4]In Western Cultures this is often referred to as Ro-Sham-Bo or Rock-Paper-Scissors

set of preferences is $P = \{pPr, rPs, sPp\}$).

In this case, if transitivity is applied to any two ordered pairs of elements in the set of preferences a contradiction with the third ordered pair arises (i.e. If transitivity is assumed on pPr and rPs then it is necessarily to have that pPs (which is clearly false)).

—

Transitivity, however, is one of many properties that *can* be possessed by a *Preference Relation.* Other such properties are *Connectivity* and *Negative Transitivity* - both of which are defined below:

1.8 Definition: Connectivity
A *Binary Relation* is said to be *Connected* if given any $x, y \in A$ with $x \neq y$ either $(x, y) \in P$ or $(y, x) \in P$.

1.9 Definition: Negatively Transitive
A *Binary Relation* is said to be *Negatively Transitive* if given any $x, y, z \in A$ such that $(x, y) \in P^c$ and $(y, z) \in P^c$ then $(x, z) \in P^c$.

Once again, to summarise these properties in a contextual fashion:
Connectivity: If x and y are different then a preference exists.
Negative Transitivity: If $x \not\succ y$ and $y \not\succ z$ then $x \not\succ z$.

Unsurprisingly, it is taken as an intrinsic property that a *Preference Relation* is a specifically tailored *Order Relation.* In turn, this enables the application of meaningful labels to describe orders with specific sets of these properties. The 3 most relevant ones are included below:

1.10 Definition: Partial Order
A *Binary Relation* is said to be a *Partial Order* if it is both irreflexive and transitive. The set of partial order relations is denoted by $\mathcal{PO}$.

1.11 Definition: Weak Order
A *Binary Relation* is said to be a *Weak Order* if it is assymetric and negatively transitive. The set of weak order relations is denoted by $\mathcal{WO}$.

1.12 Definition: Linear Order
A *Binary Relation* is said to be a *Linear Order* if it is a connected *Weak Order.* The set of linear order relations is denoted by $\mathcal{LO}$.

It is important to note that these categories of functions actually intercept. In fact, the criteria to be a *Linear Order* are more stringent than those of a *Weak Order*, which, in turn, are stricter than those required for a *Partial Order.* This gives rise to the following lemma:

1.13 Lemma *Relative Strength of Order Relations*
$\mathcal{LO} \subset \mathcal{WO} \subset \mathcal{PO}$.

Proof
The first statement here ($\mathcal{LO} \subset \mathcal{WO}$) comes straight from the definition, so all that is required is to show that $\mathcal{WO} \subset \mathcal{PO}$.

Let us assume that P is a weak order that is not irreflexive. In this case, xPx for all $x \in A$ but asymmetry means that xP^cx for all $x \in A$. This is a contradiction, hence any weak order is also irreflexive.

Now assume that P is not transitive. This is to say that; given some $x, y, z \in A$, xP^cy,, yP^cz and xPz. As P is a weak order (and so assymetric): $xPy \Rightarrow yP^cx$ and $yPz \Rightarrow zP^cy$. Given zP^cy and yP^cx, negative transitivity implies that zP^cx, which, together with the previous assumption that xP^cz and the property of irreflexivity yields a contradiction. Hence any irreflexive weak order must also be transitive, and so is a partial order.
—

The following example is posed to illustrate the differences between these types of ordering:

1.14 Example *(Differences between types of order)*
Consider the set of possible actions $X = \{a, b, c, d, e\}$ with the relationships that: $a \succ b$, $b \succ c$, $b \succ d$, $c \succ e$. Depending on the type of order in use, these preferences can be presented graphically in one of the following ways:

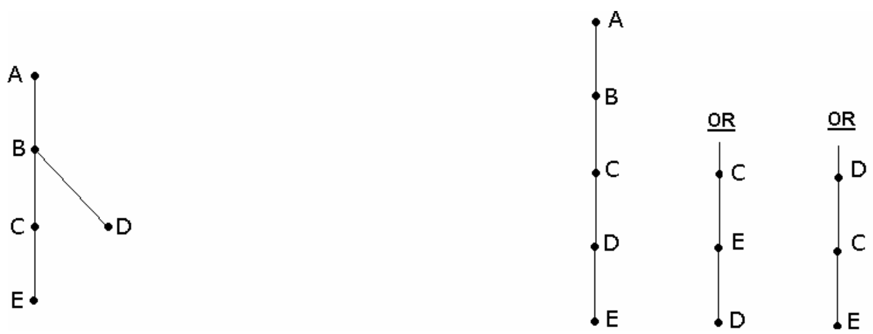

Figure 1: Left - A Partial or Weak Order. Right - Possible Linear Orders.

From the diagrams it can be seen that although there is a unique representation for a *Weak Order* or a *Partial Order*, connectivity means that this is not the case for a *Linear Order*. The simplest explanation of this is that the information given is insufficient to extrapolate preference between the 3 variations pictured. If, however, the additional relationship that $d \succ c$ (or that $e \succ d$) were included then the associated *Linear Order* would have a unique representation.

—

Simply because there would be far fewer contingencies and extraneous cases to consider, in an ideal world all orders would be complete and linear. However, this would sacrifice realism and applicability as a consequence. To pose an example, consider a choice between two different drinks (say Coke and Lemonade). In this case it is fair to say that a preference is likely to exist. However, given the same choice between two almost identical beverages (say Coke and Pepsi) it is the likelihood that no preference exists (or, at the very least, that the difference between the two choices is considered negligible) is far greater. As a consequence it could be argued that a *Partial Order* has far more "mundane realism", and is far more "rational", that its linear counterpart - if only because mankind possesses the flaw of indecision! Rationality is defined mathematically in the following way:

Definition 1.15: Rational
A preference relation, P, is said to be *Rational* if the following criteria are met:
(1) P is Connected.
(2) P is Transitive.
(3) P is Continuous.
(*If* $\bar{a} \succeq a_i$ and $a_i \succeq \underline{a}$ $(\forall\ a_i \in A)$ then $\exists t \in [0, 1]$ such that $a_i \sim t\bar{a} + (1 - t)\underline{a}$

The 3 criteria forming this definition are based largely on intuition. Connectedness and transitivity have been discussed a number of times throughout this section, whilst continuity gives us the interesting property that any action in our set can be considered in different to some "lottery" between the most and least preferable actions[6].

The subjectivity of preference finally comes to light here, as the values assigned for continuity will depend entirely on the decision maker. In addition, this value will be based purely on personal opinions of each action's value making it rather inconsistent. It is the general consensus, however, that its inclusion is necessary to model rationality closely even if it's aforementioned subjectivity causes makes generalisation unattainable.

[6]This is discussed more rigorously later

2. From Preference to Choice

At the beginning of the previous section it was stated that preference (in a mathematical sense) is a weaker condition than choice, and it is because of this fact that a model of choice can be built from the theory surrounding it. Psychologically speaking, choosing is essentially the process of culminating all the information available on (personal) preference before following it to its logical conclusion. This concept can be seen diagramatically in Example 1.14, where, in both cases, action A should be chosen.

With this in mind, a choice function can be formally defined in the following way:

2.1 Definition: Choice Function
Let X be a set of possible actions, and let $2X$ denote the set of all subsets of X. A function $C : 2X \to 2X$ is a *Choice Function* if $C(A) \subseteq A$ for every subset A of X.
Moreover, given a preference relation, P, its *Associated Choice Function* satisfies: $C(A) = \{y \in A | \nexists x \in A : xPy\}$

In essence, the *Associated Choice Function* performs a pairwise comparison over a specified subset of actions and returns the set of those which are not considered inferior to any other given action. As rationality cannot be ensured, the resultant choice set does not necessarily have a single element. This is to say that it could also be the empty set (meaning there is no clear choice) or that it contains multiple elements (meaning there is a *Partial Order* and, again, no clear choice). These possibilities are illustrated in the following examples.

2.2 Example *(Outcomes of Choice Functions)*
Consider a set of possible actions $X = \{a, b, c\}$ and a preference relation, P, which initially gives that $a \succ b \succ c \succ a$ [6]. Consider the application of C, the choice function associated to P. As the set of relations is cyclic every action is "bested" by another action. In turn, this means that $C(X)$ is the empty set and no "optimal"[7] choice can be made from the information provided.

Now assume that (instead of $a \succ b \succ c \succ a$) $a \succ c$ and that $b \succ c$. Again apply C. In this case $C(X) = \{a, b\}$ and so a and b are equally "optimal" choices based upon the relationships available, it is important to note that, in this case, it is not necessarily true that a and b are considered indifferent.

Finally consider the case where $a \succ b \succ c$. This time $C(X) = \{a\}$ and thus there is a unique "optimal" choice.

[6] See: Example 1.7
[7] In the sense that is at least as good as any other given action

A difficult question arises from this in is whether it is possible to choose between elements of a set of choices (i.e. the set returned by the *Choice Function*). For two (or more) elements to remain in the set of choices it is clear that no direct comparison can be drawn between them, however, in reality it could be possible that one is actually "more preferable" than the other(s) [8]. Although this problem can not arise in cases where rationality holds (due to completeness), it highlights an underlying flaw of the model which affects preference relations that generate a *Partial Order*. Consequently it seems sensible to look at alternative ways to model choice in an attempt to see whether this issue can be resolved.

Arguably, the most instinctive alternative method would be to define a function that assigns a value to each possible outcome, assigning a greater value to a greater preference and so effectively forcing completeness. Such a function is called an *Ordinal Value Function*, and is defined formally in the following way:

2.3 Definition: Ordinal Value Function
Let A be a set of actions and let P be a preference relation acting on it. Then $v(.) : A \rightarrow \mathbb{R}_+$ is an *Ordinal Value Function* representing P if: $\forall\ a, b \in A$

$$v(a) \geq v(b) \Leftrightarrow a \succeq b$$

Conveniently, the existence of an *Ordinal Value Function* representing a given preference relation can be guaranteed:

2.4 Theorem: Existence of an Ordinal Value Function
Given a finite set of actions A and a complete, transitive weak order P an *Ordinal Value Function* $v(.)$ representing P on A can always be constructed. Additionally, $v(.)$ is unique up to strictly increasing transformations.

Proof:
See: French. Theorem 3.3 (pg 77) and Theorem 3.5 (pg 79)

Although the requirement of completeness means that the use of *Ordinal Value Functions* does not resolve the issue of multiple, equivalent choices, replacing a preference relation with an agreeing *Ordinal Value Function* is a step closer to being able to quantify preference. Quantification of preference may seem to be of little importance when working with only one objective (i.e. in one dimension), but, when trying to balance multiple objectives, it feels almost indispensable.

It is important to note at this point that even though each action can now be assigned a value with the property that a greater value is equivalent to greater preference, these values are still rather meaningless. No useful information, for example, is gained from adding, subtracting, or taking the expectation of these

[8] Note: This is different from the phrase "referred to" in this sense - they are not interchangable

values. *Ordinal Value Functions* can, however, be used to generate a set of functions to which basic operations can be applied to gain "meaningful" information - these are known as *Utility Functions*.

If an *Ordinal Value Function* assigns values in a meaningful way (i.e. an action allocated a value of 2 is twice as preferable as an action with a value of 1) then such a function is called a *Utility Function*. Due to its roots in preference and rationality, it is important to note that such a function is still rather subjective as different decision makers may hold different commodities to be of different values (i.e. One might prefer Coke and the other Lemonade). As previously mentioned, unlike with *Ordinal Value Functions*, when working with *Utility Functions* basic operations can be performed and conclusions drawn from them. Unfortunately, *Utility Functions* can be difficult to find[9].

There are some cases in which it can be guaranteed that a *Utility Function* representing a given rational preference relation can be found. However, in order to discuss this properly it is important that the idea of rationality of preference is revisited so that some additional concepts can be introduced and incorporated into the definition to make it more stringent. The first of these ideas is the notion of a lottery:

2.5 Definition: Simple Lottery
A *Simple Lottery* on X is a function $\lambda : X \to \mathbb{R}_+$ such that:
1. $\lambda(x) \geq 0$ for all $x \in X$.
2. $\sum_{x \in X} \lambda(x) = 1$

Notation:
The set of all Simple Lotteries over possible outcomes will be denoted by $\mathcal{L}$
A lottery wherein the outcome x_i will occur with probability p_i is denoted by $l =< p_1, x_1; p_2, x_2; ...; p_r, x_r >$
For convenience, outcomes shall be labelled in such a way so as to ensure that $x_1 \succeq x_2 \succeq ... \succeq x_r$

In the context of *Choices with Risk*, it would seem fair to use a lottery to model the set of consequences associated with an action (note that $\lambda(x)$ is simply the probability of outcome x occurring). Consequence is incredibly important to consider when working in a situation where certainty does not apply as it will often have an affect on the final decision. Likewise, as the environment in which the decision is being made cannot necessarily be guaranteed, possibilities of environments may need to be taken into account. In both cases a lottery can be used to model this uncertainty, and if a number of possible environments and their consequences need to be considered simultaneously a *Lottery of Lotteries* can be used. Formally:

[9] See Appendix A

2.6 Definition: Compound Lottery
A Compound Lottery is a function $C : L \rightarrow \mathbb{R}_+$ such that:
1. $C(l) \geq 0$ for all $l \in L$.
2. $\sum_{l \in L} \lambda(l) = 1$

To justify the use of a compound lottery, consider the process of making an investment. It could be argued that investing money is simply choosing between a number of gambles in an attempt to make as much profit as possible. Although the financial market is rather unpredictable, it is expected to act in certain ways based upon current trends and empirical data meaning that it is feasible to attach probabilities to how the market might change over the period of investment. Likewise, given data about how the market will act it is possible to attach probabilities to how this might affect a company's share price - each case of which can be modelled by a simple lottery.

In general it is assumed that a *Compound Lottery* produces an outcome after a finite number of randomisations (i.e. it is not possible to indefinitely win entries into further lotteries), in which case it is said that the lottery in question is *Finitely Compounded*. This assumption/property is incredibly useful as it allows us to reduce any compound lottery yielding entries into simple lotteries as outcomes to a simple lottery modelling the same "gamble" by "multiplying through" probabilities. Under a rational decision making process, it seems consider these two gambles to be indifferent (i.e. there is no value given to the gamble itself - only to its outcome). This leads to the following axiom:

2.7 Axiom: Reduction of Compound Lotteries
Given a compound lottery, C which gives as outcomes entries into further simple lotteries, there exists a simple lottery S such that C and S are considered to be indifferent.

To continue this segway into rationality, it is necessary to define a reference lottery:

2.8 Definition: Reference Lottery
A *Reference Lottery* is defined as: $x_1 p x_r =< p, x_1; 0, x_2; ...; 0, x_{r-1}; (1-p), x_r >$

This is essentially a lottery where the only possible outcomes which are possible are x_1 with probability p and x_r with probability $(1-p)$. It seems natural to assume that in order to have rational choice the following also needs to be assumed:

2.9 Axiom: Monotonicity

$$x_1 p x_r \succeq x_1 p' x_r \Leftrightarrow p \geq p'$$

Likewise, there is no real gain in entertaining a trivial case where no preferences exist. Hence:

2.10 Axiom: Non-triviality

$$x_1 \succ x_r$$

Finally, rationality needs to include that a lottery wherein one element is replaced for another that is considered indifferent to it should be considered indifferent from the original lottery. Formally:

2.11 Axiom: Substitutability
Let $b, c \in A$ such that $b \sim c$. Choose $l \in L$ such that $l =< ...; q, b; ... >$ and construct l' from l by substituting b with c (leaving everything else the same) - ie. $l' =< ...; q, c; ... >$. Then $l \sim l'$.

So, collating these gives:

2.12 Definition: Rational *(Revised)*
A preference relation, P, is said to be *Rational* if the following criteria are met:
(1) P is Connected.
(2) P is Transitive.
(3) P is Continuous.
(4) *'Reduction of Compound Lotteries'* holds for P
(5) P is Monotonic
(6) P is Non-Trivial
(7) P is Substitutible

Using this revised definition of what it means to have *Rational Preference* in conjunction with the idea of expected utility -

2.13 Definition: Expected Utility
Given a simple lottery $l =< p_1, x_1; p_2, x_2; ...; p_r, x_r >$ and a utility function, $u : X \to \mathbb{R}$ *Expected Utility* of l is defined to be:

$$\mathbb{E}_l(u) = \sum_{i=1}^{r} p_i u(x_i)$$

- the following theorem can be deduced:

2.14 Theorem: Existence of Utility Functions
Given a rational preference relation P over A there exists a utility function $u(.)$ on A representing P in the sense that:
(1) $x_i \succeq x_j \Leftrightarrow u(x_i) \geq u(x_j)$ $(\forall\ x_i, x_j \in A)$
(2) $l \succeq l' \Leftrightarrow \mathbb{E}_l(u) \geq \mathbb{E}'_l(u)$ (for some $l, l' \in L$)
In addition, this function is unique up to affine transformations[8].

[8]Of the form $x \to Ax + b$ for some linear transformation A

Proof:
See French; Theorem 5.1 (pg 162) and Theorem 5.2 (pg 167)

This section is concluded by briefly revisiting the idea of a *Choice Function*. If a choice function over the space of lotteries based upon expected utility is constructed then it should satisfy $C(L) = \{l \in L \mid l' \in L \Rightarrow \mathbb{E}_l(u) \geq \mathbb{E}_{l'}(u)\}$. This is known as the *Utility Maximising Choice Set* and highlights the strong connection between choosing and maximising (expected) utility.

3. Adding Attributes

So far only *Choice under Certainty'* and *'With Risk'* in a single dimension have been considered. Although it is sometimes the case that a decision is being made in such a way that a single-dimensional utility function provides significant representation, it is far more common to have a decision that needs to be made based upon multiple, often conflicting, criteria or objectives, each of which possessing its own preference ordering. With choices that are both sufficiently complicated and sufficiently important enough to warrant the application of Decision Theory it is rare that the first case will arise[9] , and so it is important that the theory already covered in this essay is extended to multiple dimensions.

Given any single objective it has been shown that a utility function can be found to represent any rational preference relation over the set of actions. From this, it seems fair to extrapolate that given n objectives n (often distinct and independent) utility functions can be found to represent the rational preferences of each one individually. This is usually represented by the vector of utility functions: $u = (u_1, ..., u_n)$. Combining this with the idea of the *Utility Maximising Choice Set* from the previous section the following is attained:

3.1 Definition: Pareto Function

$$C(L) := \{l \in L | \nexists\ l' \in L \mid \forall\ i\ \mathbb{E}_l(u_i) \geq \mathbb{E}_{l'}(u_i) \\ \text{and } \exists\ i_0 \text{ such that } \mathbb{E}_l(u_{i_0}) > \mathbb{E}_{l'}(u_{i_0})\}$$

Although somewhat overly simplified, the pareto function (and the pareto rule it is based upon) provide a good starting point from which to build. A pareto function works by eliminating any lottery that is strictly inferior to any other lottery (i.e. less good or equal in all respects except one, where it is strictly worse) much like the more well-known concept of the 'process of elimination'. Formally, it is said that the *Pareto Rule* eliminates all actions that are (strictly) dominated by another:

3.2 Definition: Dominance
Given two points $p, q \in \mathbb{R}^n$, p dominates q if every component of p is at least as large as the corresponding component of q. (i.e. $p_i \geq q_i\ \forall\ i \in \{1, 2, ..., n-1, n\}$)

The problem of choice optimisation can be resolved rather simply in the trivial case where one criterion is considered to be of sufficiently greater importance to the rest that the other utilities can be ignored. In this case the *Preference Marginal* is used to make the choice:

3.3 Definition: Preference Marginal
Given two lotteries $l, l' \in L$, $l \succeq_i l'$ if $\mathbb{E}_l(u_i) \geq \mathbb{E}_{l'}(u_i)$. The preference relation $\succeq_i$ is known as the i^{th} preference marginal.

[9] This is purely conjecture on my part. My assumption is that most people would agree that the application of Decision Theory is overkill in simpler cases.

Triviality and simplification aside, the problem of "balancing" multiple objectives can prove to be rather difficult. Although it can be guaranteed that a utility function is "meaningful", there is no way of telling whether there is a legitimate reason to claim that the overall utility of a given utility vector can be calculated from its components via the process of addition, multiplication, or even at all. If a utility vector cannot be mapped in a "meaningful" fashion into $\mathbb{R}$ then is it even possible to compare the multi-attribute utilities of two (or more) actions?

The answer to this, like most things in mathematics, is largely dependent upon the properties possessed by the utility vector (and its corresponding utility functions) in question. As will now be explored independence, and different types thereof, play a key role in the separation of multi-attribute utility functions, with some concepts becoming less problematic once beyond the 2-dimensional case. For the sake of maintaining a logical order independence in two dimensions will be discussed first and then progressed from to generalise the theory .

There are a number of ways of definitions of independence that can be applied to a multi-attribute utility functions. Arguably the most simple of these is

3.4.1 Definition: Utility Independence
X is said to be *Utility Independent* of Y if; given a fixed $y_0 \in Y$, $(x, y_0) \succeq (x', y_0)$ $\Leftrightarrow x \succeq x'$ for all $x, x' \in X$.
If X is independent of Y and Y is independent of X then it is said that they are *Mutually Utility Independent*.

which states that; given lotteries with a fixed, common, level of Y and a varying level of X, preferences depend only on X. It is important to note that utility independence is not a reflexive condition. To cite a real life example consider balancing the objectives of "optimising available space" and "increasing employment". In this case it is fair to say that the number of employees would have an affect on how space is optimised/allocated, but it seems somewhat insane to say that people would be employed (or not) in order to "optimise space". The idea of independence being mutual is rather important, especially as it enables the derivation of additional properties of multi-attribute utility functions when combined with an second type of independence:

3.5.1 Definition: Additive Independence
Two objectives/attributes X and Y are *Additively independent* if $\forall\ x, x' \in X$ and $\forall\ y, y' \in Y$:

$$< \frac{1}{2}, (x, y); \frac{1}{2}, (x', y') > \sim < \frac{1}{2}, (x, y'); \frac{1}{2}, (x', y) >$$

The reasoning behind the name *Additive Independence* should be made clear from the following:

3.6 Theorem: *Additive Utility Functions*
Suppose that $u(x, y)$ is a utility function on $X \times Y$ defined in such a way that $u(x_0, y_0) = 0$. Then $u(x, y) = u(x, y_0) + u(x_0, y) + ku(x, y_0)u(x_0, y)$ for some $k \in \mathbb{R}$ iff X and Y are mutually utility independent.
Furthermore, $k = 0$ iff X and Y are *Additively Independent.*

Proof:
See French. Lemma 5.5 - Corollary 5.7 (pg 184-186) for first statement and for additive independence implying that $k = 0$

($k = 0 \Rightarrow$*Additive Independence*)
Assume $k = 0$, then $u(x, y) = u(x_0, y) + u(x, y_0)$.
Then:

$$\begin{aligned}\tfrac{1}{2}u(x, y) + \tfrac{1}{2}u(x', y') &= \tfrac{1}{2}u(x_0, y) + \tfrac{1}{2}u(x, y_0) + \tfrac{1}{2}u(x_0, y') + \tfrac{1}{2}u(x', y_0)\\ &= \tfrac{1}{2}((u(x_0, y) + u(x', y_0)) + (u(x, y_0)(x_0, y')))\\ &= \tfrac{1}{2}u(x, y') + \tfrac{1}{2}u(x', y)\end{aligned}$$

So X and Y are Additively Independent.
—

From this it can be seen that a 2-dimensional utility function can be broken down into a simpler, more calculable form providing that mutual utility independence holds. In the first case (i.e. $k \neq 0$) it is said that the utility function can be written in *Multiplicative Form* (i.e. $u(x, y_0) + u(x_0, y) + ku(x, y_0)u(x_0, y)$), whilst in the second case (i.e. additive independence and $k = 0$) it is said that the utility function can be written in *Additive Form* (i.e. $u(x, y) = u(x, y_0) + u(x_0, y)$). Generally speaking, in the multiplicative case, the easiest way to calculate the constant k is to ask the decision maker for an indifference relation between two actions[10] (i.e. of the form $(x_i, y_i) \sim (x_j, y_j)$) then inputting them into the equation and solving the two resulting simultaneous equations (this, however, requires that $u(x_i, y_0)u(x_0, y_i) \neq u(x_0, y_j)u(x_j, y_0)$). An example of this process is as follows:

[10]This process is similar to asking for initial conditions when solving differential equations

3.7 Example: *Solving a Choice Problem with Two Objectives*[11]

Problem:
A Decision Maker has decided to go on a trip to a garage in order to purchase a new car. They have decided to rank cars based upon their preferences for two attributes: 'Fuel Economy' (x) and 'Comfort' (y). The decision maker has previously determined that their preferences for x and y are both mutually utility independent, and that their respective marginal utilities are linear (with greater values preferred in both cases). In addition, they have established that they hold the following indifferences: $(41, 17) \sim (51, 12)$ and $(56, 13) \sim (46, 16)$.

The cars available in their price range when they get to the garage are; a Vauxhall Corsa - $(50, 14)$, a Honda Civic - $(43, 17)$, a Citroen C4 - $(49, 15)$, and a Fiat Panda - $(57, 13)$. Which car should they purchase?

Solution:
First, consider the indifference relation $(41, 17) \sim (51, 12)$.

A relation of the form $(x_0, y_1) \sim (x_1, y_0)$ is required, so it makes sense to let $x_0 = 41$, $x_1 = 51$, $y_0 = 12$ and $y_1 = 17$. Letting $u(41, 12) = 0$ and $u(51, 12) = 1$, observation (and the linearity of marginal utility) gives an equation for x: $u(x, 12) = \frac{(x-41)}{10}$. Likewise, letting $u(41, 12) = 0$ and $u(41, 17) = 1$ the same method can be applied for y giving the equation $u(41, y) = \frac{(y-12)}{5}$.

Applying 3.6 gives: $u(x, y) = \frac{(x-41)}{5} + \frac{(y-12)}{10} + k\frac{(y-12)(x-41)}{50}$

Now consider the second indifference relation - $(56, 13) \sim (46, 16)$.

This gives that $u(56, 13) = u(46, 16)$. Applying the above relation gives $(56, 13) = 3 + \frac{1}{10} + k\frac{3}{10} = \frac{31+3k}{10}$ and $(46, 16) = 1 + \frac{4}{10} + k\frac{2}{5} = \frac{14+4k}{10}$ meaning $k = 17$.

So a Utility Equation representing these preferences is: $u(x, y) = \frac{(x-41)}{5} + \frac{(y-12)}{10} + 17\frac{(y-12)(x-41)}{50}$

Applying this to each available car, the following Utilities are found:
$u(\text{Corsa}) = 6.12$, $u(\text{Civic}) = 4.3$, $u(\text{C4}) = 10.06$, $u(\text{Panda}) = 8.74$

So the decision maker should buy the Citroen C4.

—

The generalisation of this theory from the two-dimensional case to the n-dimensional case is surprisingly simple. In fact, if an additional case is ignored, the basic form is almost identical. Firstly, however, the definition of mutual utility independence needs to be revised in order to make it more general:

[11] Adapted from [2] pg. 188-189

3.4.2 Definition: Utility Independence *(Revised)*
Given a space of actions $X_1 \times X_2 \times ... \times X_q$ and an index set I, let $Y = \times_{i \in I} X_i$ and $Z = \times_{i \notin I} X_i$. Then, in the decomposition (Y, Z), the objectives combining to form Y are said to be *Utility Independent* of the objectives forming Z if (given a fixed $z \in Z$) preferences for lotteries on (Y, z) are independent of z.
Furthermore, $X_1, X_2, ..., X_q$ are *Mutually Utility independent* if Y is *Utility Independent* of Z for any decomposition (Y, Z). *(i.e. For any index set)*

The idea behind this revised definition of *Mutual Utility Independence* is that is that any way of distributing objectives into two disjoint spaces should yield two mutually utility independent objectives (in the sense of *Mutual Utility Independence in Two-Dimensions*) regardless of how the objectives have been grouped. *Mutual Utility Independence* is a rather strong condition, and allows us to generalise the multiplicative form of a Utility Function from the previous theorem (3.6).

3.8 Theorem: *General Multiplicative Form of Utility Functions*[11]
Given a utility vector, $u(.)$, on the space of actions $X_1 \times X_2 \times ... \times X_q$ such that $X_1, X_2, ..., X_q$ are mutually utility independent, the following equality holds:

$$1 + ku(x_1, x_2, ..., x_q) = \prod_{i=1}^{q} (1 + ku_i(x_i))$$

Where $u_i(x_i) := u(x_1^0, x_2^0, ..., x_{i-1}^0, x_i, x_{i+1}^0, ..., x_q^0)$ and $u(x_1^0, x_2^0, ..., x_q^0) = 0$

A weaker condition, again based upon utility independence, leads to an alternate generalisation of the multiplicative form. Instead of requiring (full) mutual utility independence as in the above, in this case it is only necessary to have that each X_i is utility independent of the space composed of all the remaining objectives. The result is called *Multi-Linear Form*:

3.9 Theorem: *Multi-Linear Form of Utility Functions*[11]
Given a utility vector $u(.)$ on the space of actions $X_1 \times X_2 \times ... \times X_q$ where each X_i is utility independent of $Y_i := X_1 \times X_2 \times ... \times X_{i-1} \times X_{i+1} \times ... \times X_{q-1} \times X_q$, the following equality holds:

$$\begin{aligned} u(x_1, x_2, ..., x_q) = {} & u_1(x_1) + u_2(x_2) + ... + u_q(x_q) \\ & + k_{1,2} u_1(x_1) u_2(x_2) + k_{1,3} u_1(x_1) u_3(x_3) + ... \\ & + k_{2,3} u_2(x_2) u_3(x_3) + ... \\ & + k_{1,2,3} u_1(x_1) u_2(x_2) u_3(x_3) + ... \\ & + k_{1,2,...,q} u_1(x_1) u_2(x_2) ... u_q(x_q) \end{aligned}$$

Where $u_i(x_i) := u(x_1^0, x_2^0, ..., x_i, ..., x_q^0)$ and $u(x_1^0, x_2^0, ..., x_q^0) = 0$

The third and final form to be considered arises from *Additive Independence.* As before, the definition needs to be revised for the more general case in order to state the theorem.

3.5.2 Definition: Additive Independence *(Revised)*
Given a space of outcomes $X_1 \times X_2 \times ... \times X_q$ and an index set I, let $Y = \times_{i \in I} X_i$ and $Z = \times_{i \notin I} X_i$. Then the decomposition (Y, Z) is *Additively Independent* if $\forall$ $y, y' \in Y$ and $\forall$ $z, z' \in Z$:

$$< \frac{1}{2}, (y, z); \frac{1}{2}, (y', z') > \sim < \frac{1}{2}, (y, z'); \frac{1}{2}, (y', z) >$$

Furthermore, it can be said that $X_1, X_2, ..., X_q$ are *Additively Independent* if Y is additively independent of Z for any decomposition (Y, Z). *(i.e. For any index set)*

Much like in the case of utility independence, this works by decomposing the space into two distinct spaces and checking additive independence on each of these decompositions using the definition established for the two-dimensional (two-objective) case. Finally, consider the (somewhat predictable) generalisation of the Additive case.

3.10 Theorem: *General Additive Form of Utility Functions*[11]
Given a utility vector, $u(.)$, on the space of outcomes $X_1 \times X_2 \times ... \times X_q$ such that $X_1, X_2, ..., X_q$ are *Additively Independent*, the following equality holds:

$$u(x_1, x_2, ..., x_q) = u(x_1) + u(x_2) + u(x_3) + ... + u(x_q)$$

Where $u_i(x_i) := u(x_1^0, x_2^0, ..., x_i, ..., x_q^0)$ and $u(x_1^0, x_2^0, ..., x_q^0) = 0$

So, providing that the objectives that the choice is being based upon are independent in some way, these formulae can be employed in order to make the process of calculating the utility of an action a lot less difficult. Obviously checking independence can still be quite time consuming, especially when the number of objectives is large, but I am lead to believe that there are ways clever ways to reducing this problem which allow fewer pairwise comparisons to be made.[12]

For more information, I highly recommend *Utility Maximization, Choice and Preference*[13] - from which a large proportion of the first section of this essay is based - as well as *Decision Theory: An Introduction to the Mathematics of Rationality*[14] - which proved itself invaluable throughout.

END

[11] See: Generalised utility independence and some applications by Fishburn and Keeney (1975) for proof. NB: This source is suggested in another text and has not been verified.
[12] The existence of such methods is hinted at in [2] where references are cited
[13] See: [1]
[14] See: [2]

APPENDIX A

Pflug and Römisch [3] suggest that the best way to solve (or at least lessen) this problem is by not actually choosing a *Utility Function*, but instead agreeing on a set of *Utility Functions* along with a preference relation utilising partial ordering. There are 4 sets typically considered, these are as follows:

1. The set of all *non-decreasing* functions (denoted $\mathcal{U}_{FSD}$)
2. The set of all *concave* functions (denoted $\mathcal{U}_{CCD}$)
3. The set of all *concave, non-decreasing* functions (denoted $\mathcal{U}_{SSD}$)
4. The set of all *convex* functions (denoted $\mathcal{U}_{CXD}$)

(Note: $U_{SSD} \subset U_{FSD}$ and $U_{SSD} \subset U_{CCD}$)

In case the reader is unfamiliar with the ideas of convexity and concavity, these are formally defined as:

Definition: Concave
A function, f, is said to be *Concave* if, for any two points x and y in its domain and $t \in [0, 1]$,

$$f(tx + (1-t)y) \leq tf(x) + (1-t)f(y)$$

Definition: Convex
A function, f, is said to be *Convex* if $-f$ is *Concave*.

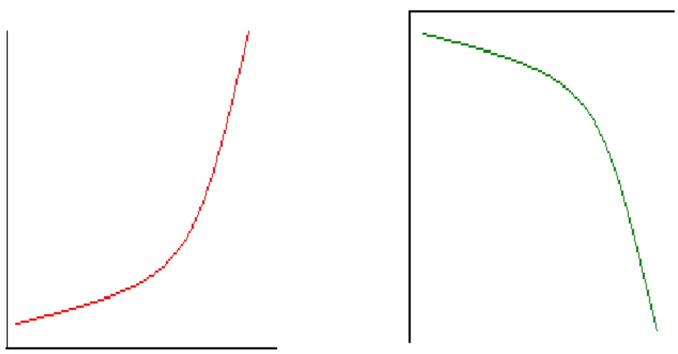

Figure 2: A Concave Function (left /red) and a Convex Function (right /green).

As far as I am aware, which of these sets of functions is used largely depends on what the model is trying to emulate.

BIBLIOGRAPHY

(1) *Utility Maximization, Choice and Preference.* Fuad Aleskerov, Denis Bouyssou and Bernard Monjardet. *Springer berlin Heidelburg.* 2007. *(Accessible online via Springerlink).*

(2) *Decision Theory: An Introduction to the Mathematics of Rationality.* Simon French. *Ellis Horwood Ltd.* 1986.

(3) *Modelling, Measuring and Managing Risk.* Georg Ch Pflug and Werner Römisch. *World Scientific.* 2007.

(4) *Aspects of the Theory of Risk-bearing.* Kenneth J. Arrow. *Suomalaisen Kirjallisuuden Kirjapaino Oy Helsinki.* 1965.

(5) *Wikipedia.* Online Encyclopedia. http://en.wikipedia.org/wiki/Main_Page

(6) *Wolfram Mathworld.* Online Mathematics Resource. http://mathworld.wolfram.com